

# Comparative study of LSA vs Word2vec embeddings in small corpora: a case study in dreams database

Edgar Altszyler <sup>\*1</sup>, Mariano Sigman<sup>2</sup> and Diego Fernández Slezak<sup>1</sup>

<sup>1</sup>Laboratorio de Inteligencia Artificial Aplicada, Depto. de Computación,  
Universidad de Buenos Aires - CONICET.

<sup>2</sup>Universidad Torcuato Di Tella - CONICET.

## Abstract

Word embeddings have been extensively studied in large text datasets. However, only a few studies analyze semantic representations of small corpora, particularly relevant in single-person text production studies. In the present paper, we compare Skip-gram and LSA capabilities in this scenario, and we test both techniques to extract relevant semantic patterns in single-series dreams reports. LSA showed better performance than Skip-gram in small size training corpus in two semantic tests. As a study case, we show that LSA can capture relevant words associations in dream reports series, even in cases of small number of dreams or low-frequency words. We propose that LSA can be used to explore words associations in dreams reports, which could bring new insight into this classic research area of psychology

## 1 Introduction

Corpus-based semantic representations (i.e. embeddings) exploits statistical properties of textual structure to embed words in a vectorial space. In this space, terms with similar meanings tend to be located close to each other. These methods rely in the idea that words with similar meanings tend to occur in similar contexts [1]. This proposition is called *distributional hypothesis* and provides a practical framework to understand and compute semantic relationship between words.

Word embeddings has been used in a wide variety of applications such as sentiment analysis [2], psychiatry [3], psychology [4, 5], philology [6], cognitive science [7] and social science [8, 9].

Latent Semantic Analysis (LSA) [10, 11, 12], is one of the most used methods for word meaning representation. LSA takes as input a training corpus, i.e. a collection of documents. A word by document co-occurrence matrix is constructed. Typically, normalization is applied to reduce the weight of uninformative high-frequency words in the words-documents matrix [13]. Finally, a dimensionality reduction is implemented by a truncated *Singular Value Decomposition*, SVD, which projects every word in a subspace

---

<sup>\*</sup>Corresponding author: ealtszyler@dc.uba.ar

of a predefined number of dimensions. Once the vectorial representation of words is obtained, the semantic similarity between two terms is typically computed by the cosine of the angle between them.

More recently, neural-network language embeddings have received an increasing attention [14, 15], leaving aside classical word representation methods such as LSA. In particular, Word2vec models [15, 16] have become especially popular in embeddings generation.

Word2vec consists of two neural network language models, Continuous Bag of Words (CBOW) and Skip-gram. In both models, a window of predefined length is moved along the corpus, and in each step the network is trained with the words inside the window. Whereas the CBOW model is trained to predict the word in the center of the window based on the surrounding words, the Skip-gram model is trained to predict the contexts based on the central word. Once the neural network has been trained, the learned linear transformation in the hidden layer is taken as the word representation. In the present paper, we use Skip-gram model, which shows better performance in [16] semantic task.

An intrinsic difference between LSA and Word2vec is that while LSA is a counter-based model, Word2vec is a prediction-based model. Although prediction-based models have strongly increased in popularity, it is not clear whether they outperform classical counter-based models [17, 18, 19].

In particular, Word2vec methods have a distinct advantage in handling large datasets, since they do not consume as much memory as some classic methods like LSA and, as part of the Big Data revolution, Word2vec has been trained with large datasets of about billions of tokens. However, often in several problems of natural and social sciences one has to form semantic embeddings based on scarce data. For example, when analyzing the semantic map of a psychiatric patient or tracking the semantic network growth in children’s writing. Moreover, this kind of approach is also relevant in sociological and linguistic research, in which linguistic patterns in word meaning networks are tracked along the time line, and small time chunks are needed [9, 20].

Which is the best method when only small amounts of data are available? In the present paper we investigate this, on the working hypothesis that Word2vec will produce very low quality embeddings when trained with small corpus, as it is a prediction-based model and it would need lot of training data in order to fit its high number of parameters. Here we examine this hypothesis, investigating the optimality of different methods to achieve reliable semantic mappings when only medium to small corpora are available for training. In these conditions, we compare Word2vec performances with LSA in a semantic categorization test, in which the capabilities of the model to represent semantic categories (such as, drinks, countries, tools and clothes) is measured.

Then we examine the performance of these models in a real-life and challenging problem based on relatively short texts: analyzing and disambiguating the content of dreams. Dream content show gender and cultural differences, consistency over time of the dreams content, and concordance of dreaming features (such as activity and emotions) with waking-life experiences [21, 22, 23]. Also, there is evidence of change in dreams contents after drug treatment [24] and shifts in content patterns in people with psychiatric disorders [25].

Most of the newest dreams content analysis methods are based on frequency word-counting of predefined categories in dreams reports [23]. A well known limitation of this

approach is the impossibility of identifying the meaning of the counted words, which is determined by the context in which they appears. For example, the occurrence of the word *fall* in a dream report may be used in different contexts, such as, *falling* from a cliff, teeth *falling* out or *falling* sick. In this context, we will test the capabilities of LSA and Word2vec on identify patterns in the usage of words among subjects. In particular, we set to analyze the semantic neighborhood of the word *run* present in the dreams reports of the different subjects. We have chosen this word because its frequency in dreams and the variety of contexts where it can be used. For example, *run* may be associated to sports activities or with chase/escape situations, which is reported to be one of the most typical dreams [26, 27].

Here we specifically analyze the capabilities of both models to identify word associations in dreams reports. We believe that word embeddings can bring new insights in dreams content analysis. On the other hand, we claim that LSA would be more appropriate in small-size corpus and should outperform Word2Vec performance in this context.

## 2 Methods

### 2.1 Semantic representations

Both, LSA and Word2vec semantic representations were generated with the Gensim Python library [28]. In LSA implementation, a tf-idf transformation was applied before the truncated Singular Value Decomposition. LSA’s representation dimensionality were tuned in order to maximize its performance in each case. In Word2vec (Skip-gram) implementations no minimum frequency threshold were used, and the window size, the number of negative samples and the representation dimensionality were tuned to maximize the performance. All other Skip-gram parameters were set to default Gensim values.

Given a vectorial representation, the semantic similarity ( $S$ ) of two words was calculated using the cosine similarity measure between their respective vectorial representation ( $\mathbf{v}_1, \mathbf{v}_2$ ),

$$S(\mathbf{v}_1, \mathbf{v}_2) = \cos(\mathbf{v}_1, \mathbf{v}_2) = \frac{\mathbf{v}_1 \cdot \mathbf{v}_2}{\|\mathbf{v}_1\| \cdot \|\mathbf{v}_2\|} \quad (1)$$

The semantic distances between two words  $d(\mathbf{v}_1, \mathbf{v}_2)$  was calculated as 1 minus the semantic similarity ( $d(\mathbf{v}_1, \mathbf{v}_2) = 1 - S(\mathbf{v}_1, \mathbf{v}_2)$ ).

### 2.2 Semantic tests

To compare LSA and Skip-gram semantic representation quality, we perform two tests in two different corpora (TASA and UkWaC): (1) a semantic categorization test and (2) a word-pairs similarity test. For each test, we studied how the performance of LSA and Skip-gram embeddings depend on the corpus size. To do this, we take 6 nested sub-samples of the training corpora, in which documents were progressively eliminated, following [29, 30]. In both cases, the minimum sub-corpus size contains only 600 documents. When any of the test words did not appear at least once in a sub-corpus, a random document was replaced with one of the discarded ones.

### 2.2.1 Semantic categorization test

In this test we measured the capabilities of the model to represent the semantic categories [31, 29] (such as, drinks, countries, tools and clothes). The test is composed by 53 categories with 10 words each. In order to measure how well the word  $i$  is grouped vis-à-vis the other words in its semantic category we used the Silhouette Coefficients,  $s(i)$  [32],

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \quad (2)$$

where  $a(i)$  is the mean distance of word  $i$  with all other words within the same category, and  $b(i)$  is the minimum mean distance of word  $i$  to any words within another category (i.e. the mean distance to the neighboring category). In other words, Silhouette Coefficients measure how close a word is to other words within the same category compared to words of the closest category. The Silhouette Score is computed as the mean value of all Silhouette Coefficients. The score takes values between -1 and 1, higher values reporting localized categories with larger distances between categories, representing better clustering.

### 2.2.2 Word-pairs similarity test

This test measures the capabilities of the model to capture semantic similarity between concepts. We used the well established WordSim353 test collection [33], which consist of 353 word-pairs (such as Maradona-football or physics-chemistry) associated with a mean human-assigned similarity score. Each word-pair is rated on a scale ranging from 0 (highly dissimilar words) to 10 (highly similar words). The evaluation score is computed as the Spearman correlation between the human scores and the model semantic similarities.

## 2.3 Case study: Semantic association in dreams reports

In this case study, we analyze the capabilities of the models to capture semantic word associations, testing whether the models embeddings can capture the semantic neighborhood of a target word in single subject’s dream series ( a collection of dream reports written by the same person). In particular, we selected the word *run* as the target word, and we focused on the detection of its distance to escape/chase contexts. The rank distance of a given word “w” with respect to *run* was measured as the rank of “w” among the cosine similarity between *run* and all other words in the vocabulary. For example, if a word has a rank of 20, it means that, among all words in the vocabulary, it is the 20th closest word using a cosine similarity metric. Finally, we will define the rank distance of escape/chase concepts as the minimum value within the ranks of the words *escape*<sup>\*</sup> and *chase*<sup>\*</sup><sup>1</sup>.

For each dream series, two independent annotators read all the dreams in which the word *run* appears, and labeled whether they refer to an escape/chase situation or not. Escape/chase situations were defined as those in which (1) someone is being chased or is under the impression of being chased or (2) someone is escaping from a real or imaginary threat. Also, for a *run* frame to be counted as an escape/chase context, it must be

---

<sup>1</sup>The asterisk (\*) refers to a word in all its forms, i.e. *escape*<sup>\*</sup> stands for *escape*, *escapes*, *escaping* and *escaped*.

associated to a negative emotional valence, thus discarding, for instance, escape/chase situations related to games or sports. With these criteria, for each dream series the annotators calculate the fraction of times the word *run* appears in an escape/chase context, obtaining a Pearson correlation coefficient of 0.98. We use the average of the annotators measurement as the ground truth, and we will refer to this values as the *escape/chase fraction*.

We used the *escape/chase fraction* as a ground truth to test the embeddings quality. Good representations should produce low rank distance in series with high *escape/chase fraction* and high rank distance in series with low *escape/chase fraction*. Thus, not only do we expect negative correlations between the escape/chase rank distance and the ground truth, but we also expect the differences in rank distances to be large when the models are trained with low and high *escape/chase fraction*.

In order to quantify these differences, we computed the linear regression of  $\log_{10}(\text{rank distance})$  vs the *escape/chase fraction*, and we used the log-linear slope as a measurement of performance. Thus, the more negative the slope is, the better the performance. It should be noted that in this analysis the series in which the word *run* appears less than 5 times were excluded.

## 2.4 Corpora

In both test, we use as training corpora the TASA corpus [34] and a random subsample of ukWaC corpus [35]. TASA corpus is a commonly used linguistic corpus consisting of 37k educational texts with a corpus size of 5M words in its cleaned form. UkWaC consists of web pages material from .uk domain. The random subsample has 140k documents with a corpus size of 57M words in its cleaned form.

For the case study we use the Dreambank reports corpus [36, 23]. The DreamBank corpus consists of 19k dreams reports from 59 subjects, containing about 1.3M words in its cleaned form.

To clean the corpora, we performed a word tokenization, discarding punctuation marks and symbols. Then, we transformed each word to lowercase and eliminated stopwords, using the stoplist in NLTK Python package [37]. Also, all numbers were replaced with the string "NUM".

## 3 Results

### 3.1 Corpus size analysis in the clustering test

As a first step for all analyses, we carried out the Skip-gram parameter optimization for both tests (Table 1). The best scores were selected to perform the corpus size analysis. In the semantic categorization test, in the case of TASA corpus,  $\text{neg} = 15$  was chosen given that it showed slightly better performance.

To compare LSA and Skip-gram embeddings quality in small size corpora, we tested both methods in random nested subsamples of TASA and ukWaC corpus (see Figure 1). Given that the appropriate embeddings dimensions depends on the corpus size [38], for each sub-corpus, we ran the models with a wide range of dimension values

		win \ neg	5	10	15
Silhouette score	TASA	5	0.107	0.107	0.109
		10	0.110	0.117	0.119
		15	0.115	<b>0.121</b>	<b>0.121</b>
	ukWaC	5	0.150	0.151	<b>0.155</b>
		10	0.146	0.149	0.151
		15	0.141	0.145	0.145
Correlation	TASA	5	0.603	0.592	0.589
		10	0.615	0.610	0.602
		15	0.623	0.618	<b>0.626</b>
	ukWaC	5	0.643	0.633	0.638
		10	0.644	0.643	0.637
		15	<b>0.647</b>	0.640	0.642

Table 1: Skip-gram’s parameter selection. Silhouette scores for the categorization test and correlations for the WordSim353 test. In all cases the embedding dimensions were set to 100.

(7,15,25,50,100,200,400), using in each case the dimension that produces the best performance.

Figure 1 shows that Skip-gram word-knowledge acquisition rate tends to be larger than LSA’s. While Skip-gram tends to produce better embeddings than LSA when they are trained with large corpora, under training with small corpora Skip-gram performance is considerable lower than LSA’s. We believe that this behavior is grounded in the fact that Skip-grams is a prediction-based model, so it requires substantial training data in order to fit its high number of parameters.

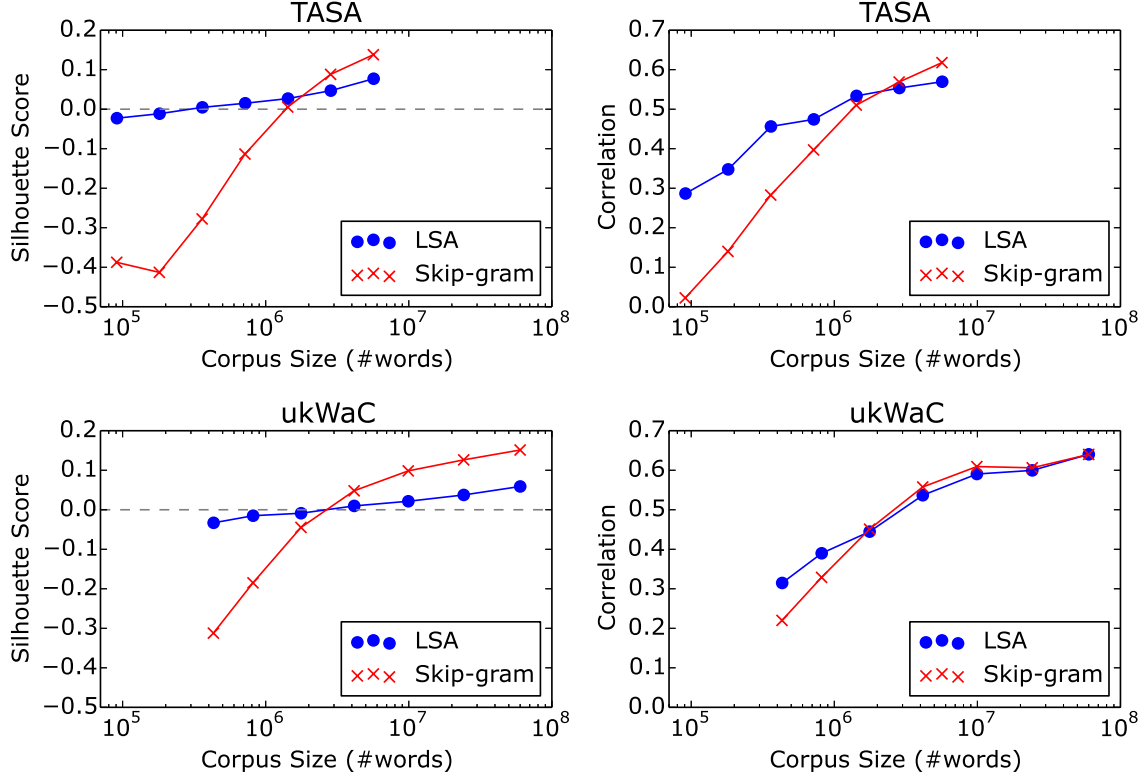


Fig. 1: Semantic categorization test performance (left graphs) and WordSim353 test performance (right graphs) in function of the corpus size for LSA and Skip-gram model. The size of the different corpus are considered as the number of tokens that they contain.

### 3.2 Case study: semantic association in dreams report

In order to check the expected differences between the associations of the word *run* in dreams and waking life, we built LSA and Skip-gram embeddings trained each corpora, and we extracted the 25 words most similar to *run* in each case. Infrequent words which appear less than 15 times were excluded. We found that word embeddings are capable of identifying differences in usage patterns of word between dreams and waking life. In TASA and ukWaC corpora, *run* is linked with words associated with a big variety of contexts, such as sports, means of transport and programming, while in dreams, *run* is directly related with words associated with chase/escape situations. For example with LSA trained in dreams we obtained words such as: *chase*, *scream*, *chasing*, *escape*, *chases*, *grab*, *screaming*, *nazi*, *hide*, *chased*, *yells*, *safety*, *devil*, *evil*, *attacking*, *killing*, *slam* and *yell*. In the same line, with Skip-gram model we found words such as: *escape*, *catch*, *chase*, *chasing*, *follow*, *dangerous*, *guards*, *robbers*, *hide*, *hiding*, *escaped*, *safely*, *safe*, *protect* and *tornado*.

Then, we tested the ability of both models to extract semantic tendencies in single dreams series following the method described in subsection 2.3. A parameter selection was made, obtaining the best performance for LSA in 200 dimensions and for Skip-gram in win=15 and neg=10 (Table 2 and Table 3).

win \ neg	5	10	15
5	-0.63	-0.96	-0.99
10	-0.95	-1.06	-1.01
15	-0.95	<b>-1.12</b>	-1.02

Table 2: Skip-gram’s parameter selection in the dreams reports analysis. The scores are the slopes in the log-linear regression of the escape/chase rank distance vs escape/chase fraction. The shown values are the mean score among 10 repetitions. The embedding dimensions were set to 50.

dim	30	50	100	200	300	400
Slope	-1.65	-1.67	-1.96	<b>-1.99</b>	-1.73	-1.77

Table 3: LSA’s parameter selection in the dreams reports analysis. The scores are the slopes in the log-linear regression of the escape/chase rank distance vs escape/chase fraction. The shown values are the mean score among 10 repetitions

The sensitivity of the method to detect both situation of the term *run* rely on the slope steepness. We expect negative correlations with steep slopes between the escape/chase rank distance and the escape/chase fraction (see Methods section for details). In figure 2 we plot the calculated distance vs the ground truth for each individual series in the selected parameters.




Fig. 2: Escape/chase rank distance vs the escape/chase fraction for each individual series. A log-linear regression was performed for one sample of LSA and Skip-gram models (blue dash-dotted line and red dashed line respectively). LSA measurements present a log-linear slope of -2.10, while the Skip-gram model has a slope of -1.11. Also, LSA and Skip-gram model show correlation of -0.57 and -0.42, with *p-values* of 0.0001 and 0.007 respectively.

While both models present a downward trend, the LSA outperforms Skip-gram with a negative log-linear slope of -1.99 and -1.12, respectively. We ran this test for 10 iterations



and slopes showed significant difference between methods (Kolmogorov-Smirnoff test,  $p < 3 \times 10^{-4}$ ).

In order to illustrate to what extent we can use these methods to explore the usage pattern of a target word in individual dream series, we show in figure 3 the 25 closest words of *run* in 3 different dreams series, using the same parameter set as in figure 2. Series 1 and 2, are the two series with the highest escape/chase fraction, while series 3 has no escape/chase situations in dreams that contain the word *run*. In the first two series, we observe that *run* neighborhood in LSA embedding contains words highly related with escape/chase situations, such as *chased* and *hide* in series 1 and *chasing*, *chases* or *trapped* in series 2. Conversely, Skip-gram embeddings do not succeed in identifying escape/chase contexts in these series. As a control case, it can be seen that series 3 do not show escape/chase related words.

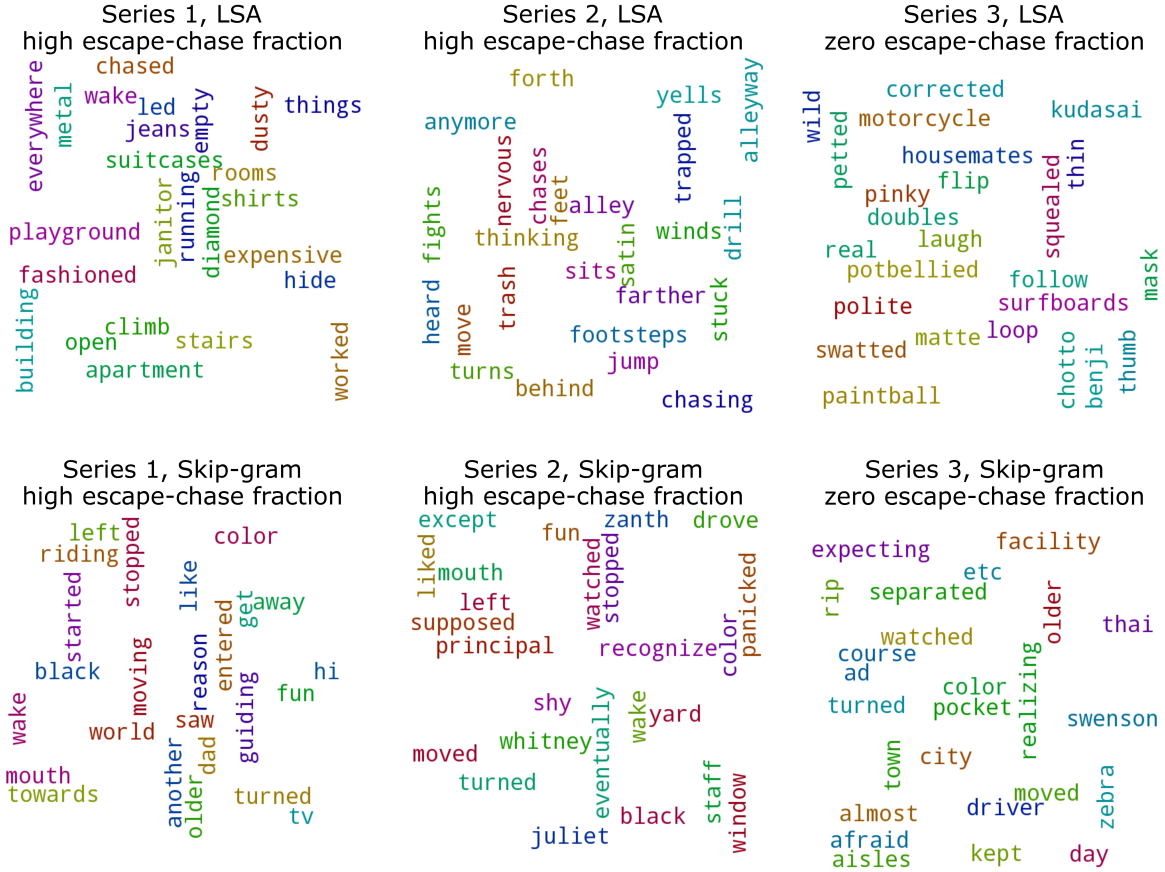


Fig. 3: Semantic neighborhood of the word *run* in 3 different dreams series for LSA and Skip-gram’s words embeddings. Series 1 and 2 are the two series with the highest escape/chase fraction, while series 3 is a control series, which has no escape/chase situations in dreams that contain the word *run*. Series 1 is the “*seventh grade girls*” series, in which only 5 of its 69 dreams reports contain the word *run* and on average 90% of these dreams refer to chase/escape situations (escape/chase fraction of 0.9). Series 2 is the “*Bay Area girls: Grades 7-9*” series, in which 6 of its 154 dreams reports contain the word *run*, 83% of which refer to a chase/escape situation (escape/chase fraction of 0.833). Series 3 is the “*Madeline3: Off-Campus*” series, in which 13 of its 348 dreams reports contain the word *run* and none of them refers to a chase/escape situation (escape/chase fraction of 0).

### 3.3 Conclusion

In the present paper, we compare the capabilities of Skip-gram and LSA to learn accurate word embeddings in small text corpora. In order to do that, we first tested the models capability to represent semantic categories (such as drinks, countries, tools or clothes) in nested subsamples of a medium size corpus. We found that Word2vec embeddings outperform LSA’s when the models are trained with medium size datasets ( $\sim 10$  millions of words). However, when the corpus size is reduced, Word2vec performance has a severe decrease, thus LSA becoming the more suitable tool. This finding gives a new insight into the prediction-based vs counter-based models discussion [17, 18, 19]. We believe that

Word2vec performance decrease in small corpora is grounded in the fact that prediction-based models need a lot of training data in order to fit their high number of parameters.

As a case study, we have studied LSA and Skip-gram capabilities to extract relevant semantic words associations in dreams reports. We found that LSA can accurately capture semantic words relations even in cases of series with low number of dreams and low frequency of target words. This is a step forward to the application of word embeddings to the analysis of dreams content. This research field addresses questions such as “what do we dream about?” and “how do gender, cultural background and waking life experiences shape the dreams content?” [21, 25, 22, 23]. We propose that LSA can be used to explore words associations in dreams reports, which could bring new insight into this old research area of psychology.

## Acknowledgments

We want to thank the teams behind the TASA [34], WaCky [35] and Dreambank [23] projects for providing us the corpora and Eduardo Schmidt for helpful discussions.

## Conflict of Interest Statement

The authors declare that there is no conflict of interest regarding the publication of this paper.

## References

- [1] Z. Harris. Word Distributional structure. 23(10):146–162, 1954.
- [2] Richard Socher, Brody Huval, Christopher D Manning, and Andrew Y Ng. Semantic Compositionality through Recursive Matrix-Vector Spaces. *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, (Mv):1201–1211, 2012.
- [3] Gillinder Bedi, Facundo Carrillo, Guillermo A. Cecchi, Diego Fernández Slezak, Mariano Sigman, Natália B. Mota, Sidarta Ribeiro, Daniel C. Javitt, Mauro Copelli, and Cheryl M. Corcoran. Automated analysis of free speech predicts psychosis onset in high-risk youths. *npj Schizophrenia*, 1(May), 2015.
- [4] Eyal Sagi, Daniel Diermeier, and Stefan Kaufmann. Identifying Issue Frames in Text. *PLoS ONE*, 8(7):1–9, 2013.
- [5] Martin Elias Costa, Flavia Bonomo, and Mariano Sigman. Scale-invariant transition probabilities in free word association trajectories. *Frontiers in integrative neuroscience*, 3:19, 2009.
- [6] Carlos G. Diuk, D. Fernandez Slezak, I. Raskovsky, M. Sigman, and G. a. Cecchi. A quantitative philology of introspection. *Frontiers in Integrative Neuroscience*, 6(September):1–12, 2012.

- [7] Thomas K Landauer. Lsa as a theory of meaning. *Handbook of latent semantic analysis*, pages 3–34, 2007.
- [8] Facundo Carrillo, Guillermo A Cecchi, Mariano Sigman, and Latent Semantic Analysis. Fast Distributed Dynamics of Semantic Networks via Social Media (Supplemental Material). 2015, 2015.
- [9] Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. Statistically significant detection of linguistic change. *Proceedings of the 24th international conference on World Wide Web (WWW '15)*, pages 625–635, 2015.
- [10] Scott Deerwester, Susan T Dumais, Thomas Landauer, George Furnas, and Richard Harshman. Indexing by latent semantic analysis. *JAsIs*, 41(6), 1990.
- [11] Thomas K. Landauer and Susan T. Dumais. A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211–240, 1997.
- [12] X Hu, Z Cai, P Wiemer-Hastings, a Graesser, and D McNamara. Strengths, limitations, and extensions of LSA. *Handbook of Latent Semantic Analysis*, pages 401–426, 2007.
- [13] Susan Dumais. Improving the retrieval of information from external sources. *Behavior Research Methods, Instruments, & Computers*, 23(2):229–236, 1991.
- [14] Ronan Collobert and Jason Weston. A Unified Architecture for Natural Language Processing : Deep Neural Networks with Multitask Learning. *Architecture*, 20(1):160–167, 2008.
- [15] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. *Nips*, pages 1–9, 2013.
- [16] Tomas Mikolov, Greg Corrado, Kai Chen, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. *Proceedings of the International Conference on Learning Representations (ICLR 2013)*, pages 1–12, 2013.
- [17] Marco Baroni, Georgiana Dinu, and German Kruszewski. Don’t count , predict ! A systematic comparison of context-counting vs . context-predicting semantic vectors. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics.*, pages 238–247, 2014.
- [18] Omer Levy and Yoav Goldberg. Neural Word Embedding as Implicit Matrix Factorization. *Advances in Neural Information Processing Systems (NIPS)*, pages 2177–2185, 2014.
- [19] Omer Levy, Yoav Goldberg, and Ido Dagan. Improving Distributional Similarity with Lessons Learned from Word Embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225, 2015.

- [20] Dionisios N Sotiropoulos, Demitrios E. Pournarakis, and George M Giaglis. Semantically aware time evolution tracking of communities in co-authorship networks. *Proceedings of the 19th Panhellenic Conference on Informatics - PCI '15*, pages 97–102, 2015.
- [21] Alan Paul Bell and Calvin Springer Hall. *The personality of a child molester: An analysis of dreams*. Transaction Publishers, 2011.
- [22] G. W. Domhoff. *Using content analysis to study dreams: applications and implications for the humanities*. Bulkeley (Ed.), New York: Palgrave., 2002.
- [23] G. William Domhoff and Adam Schneider. Studying dream content using the archive and search engine on DreamBank.net. *Consciousness and Cognition*, 17(4):1238–1247, 2008.
- [24] Nili T Kirschner. Medication and dreams: Changes in dream content after drug treatment. *Dreaming*, 9(2-3):195, 1999.
- [25] G William Domhoff. Methods and measures for the study of dream content. *Principles and practices of sleep medicine*, 3:463–471, 2000.
- [26] Tore A. Nielsen, Antonio L. Zadra, Valerie Simard, Sebastien Saucier, Philippe Stenstrom, Carlyle Smith, and Don Kuiken. The typical dreams of Canadian University students. *Dreaming*, 13(4):211–234, 2003.
- [27] Rm Griffith, O Miyagi, and A Tago. The universality of typical dreams: Japanese vs. Americans. *American Anthropologist*, 60:1173–1179, 1958.
- [28] Radim Rehuek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, 5 2010. ELRA.
- [29] John A Bullinaria and Joseph P Levy. Extracting semantic representations from word co-occurrence statistics: a computational study. *Behavior research methods*, 39(3):510–526, 2007.
- [30] John A Bullinaria and Joseph P Levy. Extracting semantic representations from word co-occurrence statistics: stop-lists, stemming, and svd. *Behavior research methods*, 44(3):890–907, 2012.
- [31] Malti Patel, John A. Bullinaria, and Joseph P Levy. Extracting semantic representations from large text corpora. *Proceedings of the 4th Neural Computation and Psychology Workshop*, pages 199–212, 1997.
- [32] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.

- [33] Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM, 2001.
- [34] S. Zeno, S. Ivens, and R. Millard, R. and Duvvuri. *The educator's word frequency guide*. Brewster, 1995.
- [35] Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, 43(3):209–226, 2009.
- [36] A. Schneider and G. William Domhoff. Dreambank. <http://www.dreambank.net/>, last accessed: Sep. 12, 2016, 2016.
- [37] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python*. " O'Reilly Media, Inc.", 2009.
- [38] Jorge Fernandes, Andreia Artifice, and Manuel J Fonseca. Automatic estimation of the lsa dimension. In *KDIR*, pages 309–313, 2011.